Chapter Three

Architectures and Control of Networked Robotic Systems

Nikolaus Correll, Daniela Rus

Networked robot systems are ensembles of robots that enhance their individual capabilities by sharing perception, computation and actuation capabilities with each other to solve problems that an individual robot could not solve alone. In this chapter, we focus on architectures and control of autonomous networked robot systems that communicate wirelessly. We will first provide an overview over current networked robot platforms spanning 3 orders of magnitude in size (from millimeter to meters in size), operation on the ground, in the air and under water, and that are networked using light, sound and radio. We will then describe classes of algorithms and their analysis used for coordination of teams of robots focusing on reactive and deliberative algorithms for sharing perception, computation, and actuation. The chapter is concluded with a summary of current challenges and promising directions.

3.1 Introduction

> A **networked robot system** is a system comprised of multiple robots in which robots actively communicate with each other, sensors, or other computational agents using some form of wireless communication.

This chapter provides an overview on technical properties of wireless communication, algorithms for networked robot systems, and their modeling and analysis. Rather than being a survey over the state of the art in networked robots, it aims at illustrating its key concepts using selected representative platforms, algorithms, and references. This chapter provides an overview over the properties of light, sound and radio communication, mechanisms for mutual localization using sound, radio and light, and an overview over reactive, deliberative and hybrid control of networked robot systems and their analysis.

Networked robot systems consist of teams of ground [Howard *et al.* (2006)], air (Chapter **??**), under water vehicles (Chapter **??**), or combinations thereof [Hsieh *et al.* (2008a)], and heterogeneous systems consisting of robots and sensor networks [Correll *et al.* (2009a)]. Networked robots are analogous to sensor networks [Estrin *et al.* (1999)], but allow individual nodes to be mobile and have manipulation capabilities. The resulting robot network is thus able to adapt the spatial distribution of its sensors and actuators, and actively modify its environment, leading to additional challenges in coordination and control. On the hardware side, robot networks not only require the capabilities for actuation and locomotion, but also sensors that are able to determine the mutual position of the robots in dynamic environments.

Name:	iSwarm	Alice + Mote	E-Puck	Khepera
Width:	0.2 cm	2 cm	7 cm	12cm
Network:	Infrared	ZigBee	Bluetooth	WiFi
Data rate:	1 Bit/s	250 kBps	1 MBps	54 MBps

Figure 3.1. Instances of networked platforms for table-top/lab experiments ranging from systems-on-a-chip of a few millimetres and communicating via infra-red to the smallest commercially available robot running embedded Linux and supporting IEEE 802.11g. From left to right: The I-Swarm robot (Chap. **??**), the Alice robot [Caprari and Siegwart (2005)], the e-puck [Mondada *et al.* (2009)], and the Khepera III (K-Team S.A.).

Instances of ground, underwater and aerial platforms are shown in Figs. 3.1., 3.2. and Fig. 3.3.. We selected these platforms as they are representative for specific communication infrastructure (light, acoustics, low and high bandwidth radio) and the order of magnitude of volume currently needed for their implementation. For instance low-range, low-bandwidth communication has been demonstrated with robots as small as 2 mm x 2 mmm (see also Chap. **??**), whereas one of the smallest platforms that implements IEEE 802.11g radio communication and a Linux networking stack (Khepera III[1]) has a diameter of 12 cm. Similarly, long-wave radio communication requires a certain minimum size due to the necessary length of the antenna [Kottege and Zimmer (2008)].

Applications for networked robot systems are generally focussing on increasing the spatial resolution of sensing and actuating robots by mobility. Applications such as assembly of structures [Yun *et al.* (2009)], environmental monitoring [Choi *et al.* (2010)], inspection [Correll and Martinoli (2009)] and deployment of wireless networks [Correll *et al.* (2009b)] directly benefit from the spatial distribution of the networked system and the ability to conduct these tasks in parallel. Applica-

[1] http://www.k-team.com

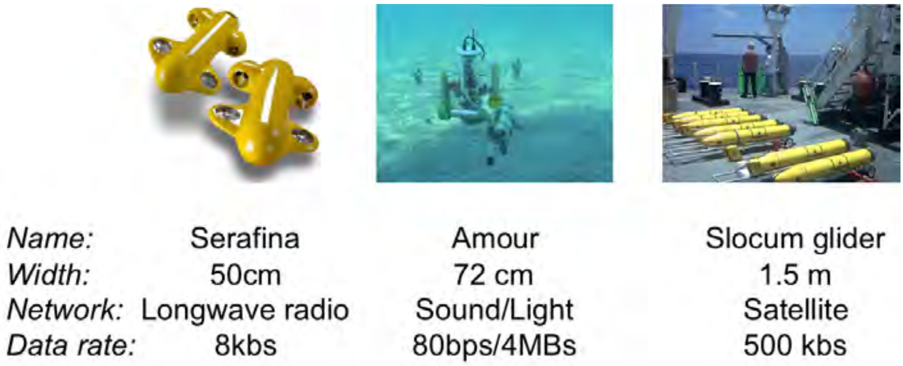

Figure 3.2. Instances of networked platforms for underwater operation that use sound, radio and light for peer-to-peer communication. Serafina [Kottege and Zimmer (2008)], Amour [Vasilescu *et al.* (2005); Doniec *et al.* (2009)] and Slocum [Rudnick *et al.* (2004)]. The Slocum glider is networked with other robots using a centralized server with which it communicates when surfacing.

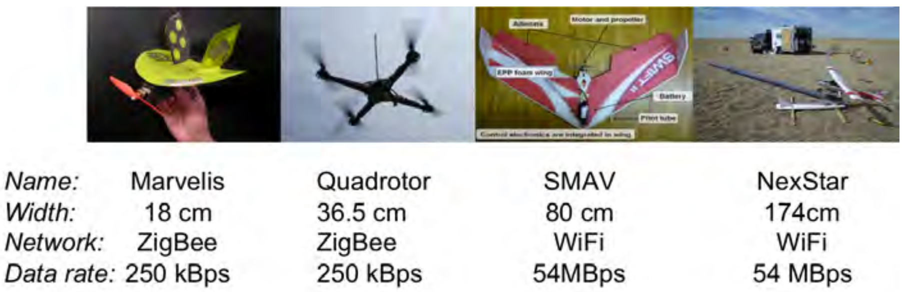

Figure 3.3. Instances of networked platforms for aerial operation that communicate using ZigBee and IEEE 802.11g. Marvelis (Mohseni lab, CU Boulder), the UFO [Gurdan *et al.* (2006)], the SMAV [Hauert *et al.* (2009)], and the NexStart UA [Frew and Brown (2008)].

Figure 3.4. Instances of networked robotics systems. From left to right: The Distributed Robot Garden at MIT CSAIL [Correll *et al.* (2009a)]. The automated warehouse robots from KIVA systems, image courtesy of Raffaelo D'Andrea. Robots interacting with visitors and people tracking systems in a shopping mall [Shiomi *et al.* (2009)].

tions such as search [Lochmatter and Martinoli (2009)] and mapping [Howard *et al.* (2006)] benefit from the increased resolution of the robot system. Finally, applications such as data ferrying [Bhadauria and Isler (2009)] benefit from the increased mobility of the networked robot system. Some of these applications are depicted in Fig. 3.4. showing robots to coordinate tending of plants [Correll *et al.* (2009a)], performing pick-up and delivery tasks in a warehouse and guiding people in a shopping mall [Shiomi *et al.* (2009)]. Soccer-playing robots, which exemplify tightly coordinated distributed sensing and actuation are treated in Chapter **??**.

3.2 Architecture: Communication and Localization

This section provides an overview over different communication technologies commonly employed in networked robot systems including communication by radio, light and sound as well as their fundamental properties. We then describe mechanisms, algorithms and systems for mutual localization, which are an important component of every networked robot system as it allows them to reason and coordinate spatially.

3.2.1 *Communication*

Two robots can communicate if they both are within a certain distance of each other, leading to a disc model for communication in its most abstract form. This model emerges from the assumption that the power of a signal that is traveling from location p_i to position p_j fades exponentially with distance, approximately following an Eq. [Chen and Kobayashi (2002)] of the form

$$P(\|p_i - p_j\|)[dBm] = P(r_0)[dBm] - 10\alpha \log\left(\frac{\|p_i - p_j\|}{r_0}\right) \tag{3.1}$$

where $P(r_0)$ is the reference power at distance r_0 from the emitter, and α is the path loss exponent which is constant and a function of the environment. While the exponent α is largely due to the so-called *path-loss* effect that corresponds to the energy lost while traveling through a medium such as air or water, real radios are also subject to *multi-path* fading, which is not reflected in Eq. 3.1. Due to the uniform emission of radio waves from the antenna and reflection of the signal from objects in the environment, the actual signal will arrive via multiple paths. The resulting interference can either be constructive or destructive, i.e. the overall amplitude is increased when waves are in phase or decreased when waves are out of phase. Thus, the multi-path effect is a *small-scale* effect that can be observed for variations in distance that are in the order of a wavelength and small variations in position of sender and receiver can drastically change the perceived signal strength. This phenomenon can also be exploited to actively improve signal strength by exploring various nearby locations, which has been demonstrated in [Lindhé *et al.* (2007)]. The combined large-scale and smale-scale effects lead to a power density distribution quite different from a disc with uniform distribution. Experimental results in quasi-ideal experimental conditions, recorded using an unmanned aerial vehicle (UAV) in a radio-free zone on a mesa in the Colorado prairie, is shown in Fig. 3.5. [Frew and Brown (2008)]. The reader is referred to [mos (2009)] for a comprehensive overview on wireless signal propagation models in a robotic's context, including models that consider large-scale fading due to obstacles in the environment.

The obtainable signal strength S together with the background noise N now

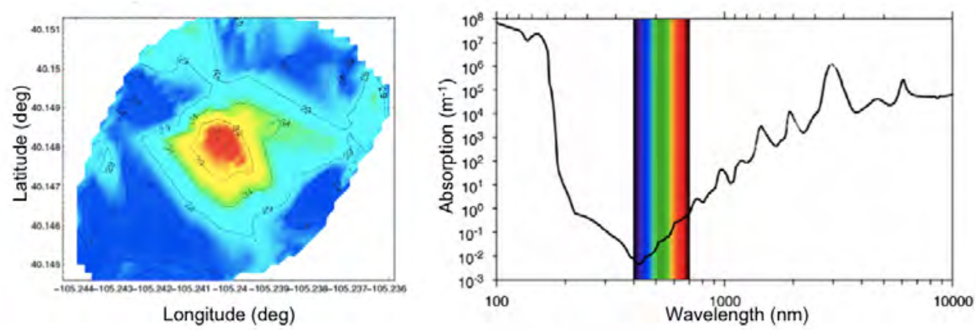

Figure 3.5. *Left:* Effective transmission power recorded by an UAV experiment in a "radio-free-zone" courtesy of Eric Frew ([Frew and Brown (2008)]). *Right:* Absorption coefficient of water from [Stomp *et al.* (2007)]. Absorption is lowest in the visual spectrum (super-imposed) and for very low frequencies with wavelengths in the order of meters (not shown). Adapted by permission from Macmillan Publishers Ltd: Nature, [Stomp *et al.* (2007)], copyright 2007.

governs the maximum throughput that can be obtained, which is known as the *Shannon-Hartley* theorem. The Shannon-Hartley theorem provides an upper bound on the capacity C of an analog channel

$$C = B\log_2(1+\frac{S}{N}) \tag{3.2}$$

with B the bandwidth of the channel in Hertz, S the total received signal power over the bandwidth measured in Watt, and N the total noise over the bandwidth measured in Watt. Notice that *bandwidth* here refers to the actual frequency band in which the signal operates. The ratio S/N is also known as the *signal-to-noise* ratio. As S decreases while N remains constant with increasing distance, the effective communication range is limited by the desired capacity of the channel. In practice modern wireless protocols therefore automatically switch data rates based on the perceived signal-to-noise ratio and thus trade-off effective range with available channel capacity. For further information on the physical principles behind path-loss and environment-specific models for radio communication the reader is referred to [Seybold (2005)].

At a higher level, the reliability of communication (whether it be radio, light or sound) is a function of the number of transceivers sharing the same channel and the chosen channel access protocol and collision mitigation mechanisms, which is beyond the scope of this chapter and the reader is referred to [Tanenbaum (2002)]. Assuring reliable communication in a multi-robot system, where transceivers often have to compete with static infrastructure communicating on the same channel is a major challenge both in robotics and in systems research.

3.2.1.1 *Radio*

Trends in radio communication in networked robot system are highly correlated with those in the personal computer industry, due to the easy availability of hardware and software as well as the low price of a mass-market product. Particularly noteworthy from this perspective are IEEE 802.11g (wireless LAN) and Bluetooth.

The IEEE 802.11g standard, which was introduced in 2003, supports a set of discrete data rates between 6 and 54 Mbit/s, which can be automatically chosen as a function of the perceived signal-to-noise ratio (see above), allowing commodity radios to communicate over distances of 300 m in open space. 802.11g radios are also available as USB sticks, Compact Flash, and SD cards, as well as as so-called system-on-a-chip that combine the radio transceiver with a powerful MIPS processor in a single package, enabling stand-alone Wifi systems using embedded Linux, e.g. OpenWRT, for a few dollars and the size of a credit card, see also [Correll *et al.* (2009b)] for a robotic system using such hardware. Using IEEE 802.11g has the advantage that standard networking protocols, such as TCP and UDP, and tools, such as mesh-networking routing protocols, are readily available, although at cost of a relatively high power consumption of 2 to 4 W.

Low-power alternatives to IEEE 802.11b are Bluetooth (IEEE 802.15.1) and ZigBee (IEEE 802.15.4 defining the physical and medium access control layer), which also operate at up to 2.4 GHz, although providing much lower data rates and ranges. Due to the large variety of Bluetooth devices and ongoing evolution of the standard, Bluetooth transceivers can provide ranges from 1 to 100 m and data transfer rates ranging from hundreds of kbps to tens of MBs. ZigBee radios operate in the range of 250 kbit/s per channel in the 2.4 GHz band, 40 kbit/s per channel in the 915 MHz band, and 20 kbit/s in the 868 MHz band. (The 2.4 GHz band is only available in Europe, whereas ZigBee transceivers for the US market operate in the 868 and 915 MHz bands.) Transmission range is between 10 and 75 meters. Open-source schematics of Bluetooth and Zigbee implementations are available for the E-puck robot[2], and ZigBee solutions have been integrated as small as 2 cm x 2 cm in [Correll and Martinoli (2009)].

A major drawback of Bluetooth for networked robot applications is its limitation to communicate with only 7 devices at a time in a master-slave hierarchy called a *piconet*, whereas mesh-networking stacks are available for ZigBee and have been used to network a swarm of quad-rotor helicopters in [Julian *et al.* (2009)].

3.2.1.2 *Communication using Light*

Communication using light, specifically in the infra-red part of the spectrum, has found wide-spread application especially in miniature robotics due to its relatively simple implementation, low cost, availability of peripherals such as remote controls, and its potential for being multiplexed as a distance sensor. For instance the miniature robot *Alice* [Caprari and Siegwart (2005)] uses its infra-red distance sensors for directed infra-red communication up to 6 cm and data rates up to 4 Bps and has the capability to receive commands from a conventional TV remote using an additional receiver mounted on its top. Although light also provides higher data rates with an appropriate coding scheme, e.g. the IrDA standard that was popular on notebooks in the 90ies supported up to 16MBit/s, light is rarely used for communication in networked robots as it requires a clear line-of-sight and is cross-sensitive to sunlight. Advantages of light-based communication are when the directional properties are actually desired, for example when used as a rela-

[2] http://www.e-puck.org

tive range and bearing device (Sec. 3.2.2), or when robots need to communicate in a medium that strongly absorbs radio waves such as water (see Fig. 3.5. for the frequency-dependent absorption properties of water). Here, blue and green light yield the best possible range [Doniec *et al.* (2009)], depending on the actual water color.

3.2.1.3 *Communication using Sound*

Sound is unlike light and radio not an electro-magnetic, but a pressure wave. Although sound is also subject to multi-path fading and attenuation, sound waves travel at only around 300 m/s as opposed to 300.000 km/s, making signal propagation time a serious concern. Also, due to the low frequency that is feasible using sound (from the audible spectrum in the order of Hz to thousands of kHz in ultrasound), the achievable data rate is intrinsically limited. Therefore, communication using sound is of little importance in networked robot systems except for speciality applications such as range and bearing (Sec. 3.2.2), human-robot interaction and underwater communication.

3.2.2 ***Mutual Localization***

Communication abilities are crucial for localization. Be it by receiving and triangulating position information from Global Positioning System Satellites or by exchanging range and bearing information with neighboring robots. This chapter focusses on mutual localization, i.e. the capability of robots to localize each other relative to each other. Mutual localization is a key capability of a networked robotic system, as it allows it to reason and communicate about spatial data. Localization using communication devices works using two fundamentally different approaches, time-of-flight measurement and fading. This section focusses on architectures for achieving relative range and bearing using sound, light and radio in water and air. In a networked robot setup mutual observations can be further refined by sharing observations within the team, which allows for improved accuracy by sensor fusion [Howard *et al.* (2003)].

3.2.2.1 *Localization using Sound*

Localization using non-audible sound, e.g. Sonar, usually relies on the time-of-flight of a signal, and can be used for both mutual localization as well as collision avoidance. As mutual localization using sound requires synchronized clocks, mutual localization using sound is usually achieved in conjunction with some other form of radio communication, which is emitted simultaneously and serves as a reference event. The actual distance between two robots is then calculated using the time-of-flight, but bearing can only be estimated by triangulating distance information from multiple neighbors [Nagpal *et al.* (2003)]. An implementation of this idea is the *Cricket* localization system [Priyantha *et al.* (2000)] that has found wide-spread use in localizing miniature robots. However, localization using sound alone is uncommon in networked robot systems, except in underwater systems due to the absorption coefficient of water discussed above (see also Chapter **??**).

Also, sound-based time-difference-of-arrival systems are prone to interferences from neighboring devices, in particular when the sound pulses are not encoded with information. For instance, in [Kottege and Zimmer (2008)] a system for underwater range and bearing estimation is presented that relies on sound emissions from a pair of projectors, which is received by a pair of hydrophones. Here, a longwave radio is used together with a scheduling algorithm [Schill and Zimmer (2007)] to ensure that only one AUV is broadcasting sound at a time in a local neighborhood. Similar to the Cricket system, the radio signal is also used as synchronization event for the time-of-flight calculation.

3.2.2.2 *Localization using Infrared*

Localization using infra-red is attractive for small-scale robot systems as it can be implemented on very small footprints (see also Chap. **??**), uses little power, and its intensity is strongly distance and direction related. For instance the Alice robot [Caprari and Siegwart (2005)], encodes its 4-bit unique identification number and the bearing of the originating transceiver into messages, which allows robots to identify a neighboring robot, estimate its distance, and calculate its relative bearing as well as the bearing of the other robot. A more capable system has been implemented for the e-puck robot, providing 5kbps communication, 1 cm in range and 2^o in bearing at distances below one meter and is available open-source [Gutierrez *et al.* (2008)].

3.2.2.3 *Localization using Radio*

The proliferation of wireless radios on robots has made using radio for mutual localization popular. Challenges of this approach are the strong spatial dependence of the perceived signal strength (see also Fig. 3.5.) as well as the fact that radio antennas employed for communication are usually omnidirectional and thus do not allow to infer the direction of a signal. While multiple techniques exist to localize a mobile robot using static infra-structure [Bahl and Padmanabhan (2000)], this chapter focusses on mutual localization of robots without relying on infrastructure in the environment. While localization techniques relying on sound and light have found widespread use, localization only using radio is still in its infancy. So far, algorithms rely solely on the large scale fading effects which lead to a reduced Radio Signal Strength Indication (RSSI) for growing distance and that can be measured using commodity hardware. Due to the strongly non-linearities of the effects underlying the perceived RSSI, distance estimates are crude and bearing estimates relying on a single measurement are impossible. In fact, multi-path fading effects can only be canceled out by actively moving around in the order of a wave-length. Only then, RSSI values can be averaged to give a crude estimate of range (see also [Lindhé *et al.* (2007)]). Dantu and Sukhatme [Dantu *et al.* (2009)] present an algorithm to estimate the bearing of a neighboring robot performing multiple measurements (in the order of hundred) on 8 locations lying on a circle with multiple meters diameter (2–5 m). By performing a principle component analysis (PCA) of the covariance matrix of these measurements, the first

Eigenvector of the covariance matrix lies along the direction with the highest variance and thus corresponds to the bearing of the other robot. Using this approach, they show measurements that are around 20^o accurate in bearing. For improving upon these results — which take multiple minutes to acquire and require the other robot to remain static — [Dantu *et al.* (2009)] suggests to exploit small-scale effects as demonstrated in [Kusy *et al.* (2007); Maroti *et al.* (2005)], which require less motion to gather. Another promising direction is to use MIMO transceivers with multiple antennas that have been originally designed to increase bandwidth by spacedomain-multiplexing (beam-forming), which is an effect that can naturally lead to an accurate bearing estimate. Similarly, localization using ultra-wind band coding is of potential interest for the networked robotics community as it is less susceptive to multi-path fading; Commercial systems that achieve 30cm accuracy in 3D over up to 160m using this technology are available from Ubisense Inc. although require at least two static sensors to pick up the signal of an active transmitter for triangulation.

3.3 Modeling and Control of Networked Robot Systems

> Coordination paradigms for networked robot systems can be classified into **reactive**, **deliberative** and **hybrid** approaches, where hybrid approaches consist of deliberative algorithms — in the simplest form implemented by simple discrete logic — that switches between different continuous, reactive dynamics.

Although all deliberative mechanisms are strictly speaking hybrid when implemented on a robotic platform involving any sort of feedback control, e.g. for motor control, a common assumption in deliberative systems is to abstract the dynamic aspects of the system away and to plan using a discrete representation of the environment.

The choice of the control architecture drastically influences the performance of the overall system as well as what can be said about a system analytically. As a rule of thumb, performance increases with the amount of *deliberation* and *coordination* within the team. At the same time, deliberative algorithms that abstract the continuous dynamics of the individual agents usually allow for making stronger guarantees on optimality and completeness. For instance, reactive systems are usually shown to converge to a local minima but cannot guarantee optimality, whereas an exhaustive search on a discrete state space can do so. The flip side, however, is that deliberation and coordination require sensing and communication abilities that might be technically infeasible in a specific domain. For instance, micro-robots such as the i-Swarm robot (Fig. 3.1., left) will require reactive coordination due to computational, sensing, and communication limitations, whereas the Linux-based Khepera robot (Fig. 3.1., right) has the ability to reason on fast amounts of data. As the performance is intrinsically limited by the quality of sensing and communication — which serves as the basis for deliberation — hardware capabilities are in fact the limiting factor on performance and finding the right trade-off between system complexity and resulting performance remains an interesting challenge.

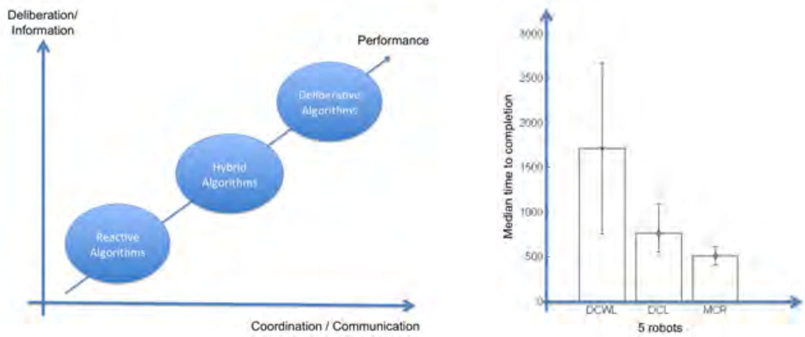

Figure 3.6. *Left:* Reactive, hybrid and deliberative approaches can be loosely classified by the amount of coordination and deliberation that they use. More coordination and deliberation will usually lead to improved performance, but comes with increased requirements on sensing, computation, and communication. *Right:* Performance (time to completion) of a distributed coverage task using a series of deliberative algorithms with increasing communication and computational complexity [Correll and Martinoli (2009)]. Whereas providing global localisation and communication drastically increases performance (from DCWL to DCL), near-optimal coordination using a market-based algorithm (MCR) provides relatively small gain for massively higher communication and computation requirements as the theoretically possible speed-up cannot be achieved with the limited sensors and actuators of the platform.

The relation between planning, coordination, and performance with respect to reactive, deliberative and hybrid coordination approaches is also show in Fig. 3.6., left. Fig. 3.6., right, illustrates the trade-off between system complexity and performance using data from an experiment in multi-robot coverage that is described in [Correll and Martinoli (2009)]. A team of 5 robots needs to cover (visit at least once) an environment that is discretized into a 5x5 grid using different deliberative algorithms. Performance increases consistently when augmenting simple individual planning (DCWL) by exchanging information on task progress within the team (DCL) [Rutishauser *et al.* (2009)], and by market-based allocation (MCR) of tasks based on a global metric [Amstutz *et al.* (2009)]. Performance improvements come at considerable cost, however: whereas exchanging information on task progress requires global localization and communication to improve upon a non-collaborative policy, up-front allocation of tasks requires orders of magnitude more computation as the allocation problem is NP-hard [Amstutz *et al.* (2009)]. In addition, the effective performance increase is well below the theoretical expectation due to limited actuator accuracy (wheel-slip) of the Alice platform. Therefore, the DCL approach might be an optimal trade-off between performance and system complexity for this particular system.

3.3.1 *Modeling and Control of Reactive Systems*

Reactive approaches are either described implicitly by heuristics that eventually lead to the desired behavior, but do not lead to analytically tractable trajectories, or explicitly as distributed control law on a global cost function, which allows to prove properties such as convergence and stability analytically.

The simplest reactive controllers implement a direct mapping between the sensors and the actuators of a robot and are known as *Braitenberg* vehicles. For this kind of controllers, sensors can also include communication devices, and they can be best formally described by artificial potential fields [Arkin (1998)] or virtual physics [Spears *et al.* (2005)]. Both approaches rely on the assumption that a reactive controller accurately tracks a vector field in the environment. The magnitude and direction of these vector fields are a function of local sensing and communication. Formally, the vector field can be described as a cost function J, leading with x_i the position of robot i to the feedback control law

$$\dot{x}_i = -\frac{dJ}{dt} \tag{3.3}$$

Different objectives, such as collision avoidance or tracking a desired shape, of a robot can be overlaid, i.e. summed up in the cost function. For example, collision avoidance between robot i and j that have distance r_{ij} is commonly (e.g., [Moshtagh *et al.* (2009)]) encoded as

$$f_{ij} = \frac{d_0}{\|r_{ij}\|} + \log \|r_{ij}| \tag{3.4}$$

where d_0 is the desired distance between the robots and resulting in the gradient

$$\nabla_{r_{ij}} f_{ij} = \frac{r_{ij}}{\|r_{ij}\|}\left(\frac{1}{\|r_{ij}\|\|} - \frac{d_0}{\|r_{ij}\|^2}\right) \tag{3.5}$$

This leads to a control law of the form

$$\dot{x}_i = -k\nabla_i \phi_i(x_i, \Gamma_i) - \sum_{j\in\Gamma_i} \nabla_i g_{ij}(q_i, q_j) \tag{3.6}$$

with $\phi_i(x_i, \Gamma_i)$ a location and communication dependent artificial potential field, k a proportional gain, and Γ_i the set of neighbors of robot i.

Using a Lyapunov-style proof to show that $\lim_{t\to\infty} \dot{x} \to 0$, often only for the components of the controller tracking the global metric of interest, one can then show that the system eventually reaches a *local* optimum of J.

For instance, [Schwager *et al.* (2009b)] derives a decentralized, adaptive control law to drive a network of mobile robots to an optimal sensing configuration (Fig. 3.7.). The control law is adaptive in that it uses sensor measurements to learn the distribution of sensory information in the environment. It is decentralized in that it requires only information local to each robot. The controller is then improved by introducing a consensus algorithm to propagate sensory information from every robot throughout the network. Convergence and consensus of parameters is proven with a Lyapunov-type proof. The controller with and without consensus is demonstrated in numerical simulations. These techniques are suggestive

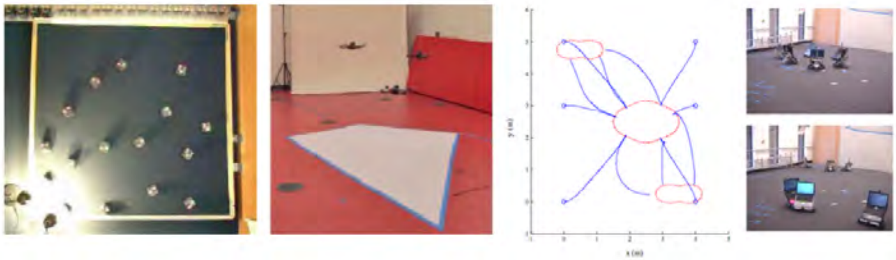

Figure 3.7. Instances of networked robots that are coordinated by purely reactive, gradient-based controllers. A robot swarm for optimal coverage of a sensory function on the ground (left, [Schwager *et al.* (2009b)]), in the air (middle, [Schwager *et al.* (2009a)]), and tracking a desired shape [Chaimowicz *et al.* (2005)]. The spatio-temporal behavior of these examples is exclusively governed by an artificial potential field and is achieved without any discrete state transitions on robot and team level.

of broader applications of adaptive control methodologies to decentralized control problems in unknown dynamic environments.

Other applications that have successfully been encoded using a global cost function with a provably convergent feedback control law are ground coverage using aerial cameras [Schwager *et al.* (2009a)] (Fig. 3.7., middle), assembly of truss structures [Yun *et al.* (2009)], tracking of global formations [Chaimowicz *et al.* (2005); Hsieh *et al.* (2008b)] (Fig. 3.7., right), and connectivity maintenance in a networked robot system, [Hsieh *et al.* (2008a)], among others. The main challenges with this approach are choosing a cost function that is convex with respected to the desired final configuration of the robotic system, i.e. lends itself to a feedback controller that controls a robot's trajectory, and actually showing its convergence. Another challenge with gradient-based approaches is the fact that usually only convergence to a local optima can be proven, but not the degree of optimality that this solution will achieve in the worst case.

3.3.2 *Modeling and Control of Hybrid Systems*

As soon as the dynamics of a robot contains discrete logic that let it switch between different continuous behaviors, we speak of a hybrid system. Early hybrid controllers are the subsumption architecture [Brooks (1986)] where various reactive behaviors subsume each other as a function of sensor input and the state of the robot. While the subsumption architecture allows for implementing complex behaviors with a minimum of computation, designing controllers is often difficult as the resulting behavior is emergent, in particular due to non-linear interaction with other robots in the team.

Given a set of reactive robot controllers, each associated with a state $q \in Q$, with Q a set of states, the robot controller can be described by Finite State Automaton (FSA) that consist of a 5-tuple $(Q, \Sigma, \delta, q_0, F)$, where

- Q is a finite set of states
- Σ is a finite set of input symbols
- $\delta : Qx\Sigma \rightarrow S$ a transition function

- $q_0 \in Q$ the initial state
- F a set of final states

Figure 3.8. Instances of networked robots with hybrid reactive-deliberative controllers. Autonomous deployment of wireless infrastructure ([Correll *et al.* (2009b)], left), distributed manipulation ([Martinoli *et al.* (2004)], middle), and inspection of regular structures ([Correll and Martinoli (2006a)], right). The spatio-temporal behavior of each of these systems is generated from the interplay of reactive controllers and discrete state transitions at the individual controller level.

For instance, in the *Stick-pulling experiment* where a swarm of robots is concerned with pulling sticks out of the ground and two robots are needed to successfully complete the task [Martinoli *et al.* (2004)] robots can be searching the arena for a stick, avoiding other robots, pulling the stick, waiting for collaboration, and completing the pull-out action in case a second robot joins them at their site. All of these are discrete states of a FSA that are associated with continuous dynamics such as randomly exploring the arena, avoiding other robots, or activating the robot's gripper. Input symbols of this system are events such as detection of a stick, detection of a robot, or detection of another robot pulling on the stick, which eventually lead to state transitions.

If the input symbols are probabilistic, this automaton reduces to a Markov chain. Common assumptions are that the likelihood that a robot encounters symbol $\sigma \in \Sigma$ is p_σ, and the likelihood of the robot to be in state q is $p_q(kT)$, with k a discrete time-step of the system with length T. This leads to the Master equation that describes the likelihood to be in state q at time kT as

$$p_q(kT+T) = p_q(kT) + \sum_{q' \in Qq} \left(p_{q'q}(kT+T)p_{q'}(kT) - p_{qq'}(kT)p_q(kT)\right) \tag{3.7}$$

with $p_{q'q}(kT+T) \in \delta$ the conditional probability to transition from state q' to state q [Correll (2007)].

Given N robots in the system, the Master equation leads itself to a rate equation with $N_q(kT)$ describing the average number of robots in state q:

$$N_q(kT) = Np_q(kT) \tag{3.8}$$

It turns out that the rate equation model (see also [Lerman *et al.* (2005)]) is a powerful tool to study the collective behavior of a hybrid networked robot system. For instance in [Martinoli *et al.* (1999)] and [Martinoli *et al.* (2004)] rate equation models have shown to provide qualitative and quantitative agreement with experimental studies in which swarms of robots aggregated pucks in an arena and collaboratively pulled sticks out of the ground, respectively. Similarly, [Correll and Martinoli (2007)] demonstrates quantitative agreement between rate equation models and a simulated swarm-robotic aggregation task, and [Correll and Martinoli (2009)] provides an accurate model of a inspection task with a swarm of up to 40 miniature robots. Other examples include aggregation of sticks [Agassounon *et al.* (2004)] or coalescence of networked robot swarms [Winfield *et al.* (2008)] where macroscopic equations provide a close match with simulation results.

Rate equation models can not only predict the population dynamics of the networked robot swarm, but can also be used to optimize its control parameters. In [Martinoli *et al.* (2004)] the rate equation model is used to estimate the optimal waiting time in the Stick-Pulling experiment and also predicts a bifurcation in the system when there are more robots than sticks (a case which supports infinite waiting time without starvation). In [Correll and Martinoli (2006c)] rate equation models are used with a dynamic optimal control framework to find optimal collaboration policies in a swarm-robotic inspection experiment and a task allocation case study [Correll (2008)].

Key challenges with this approach are stability of the rate equation systems, in particular for large numbers of possible interactions between robots — which lead to non-linear rate equations with strong coupling, and the inability of the rate equation approach to capture rare events that have significant impact on the robotic system. Also, deriving the probabilities for state transitions in the system from geometric properties of the environment and the robot controller is often infeasible. Here, system identification [Correll and Martinoli (2006b)] is a promising direction. Finally, the rate equation approach encounters difficulties with system that have rare events that lead to drastic changes in the system behavior, which are better captured by stochastic modeling techniques such as Monte-Carlo simulation. See also Section **??** of this book for an overview over rate equation-based models including case-studies on object clustering and collaborative decision-making.

A different approach to modeling and control of hybrid systems is to describe the discrete automaton and possible transitions between states using temporal logic specifications [Kloetzer and Belta (2010)]. This approach allows to separate the continuous dynamics (given by the robot's motion) of each state from the discrete dynamics of state evaluation, and allows to leverage tools from formal verification to prove properties of the system.

3.3.3 *Modeling and Control of Deliberative Systems*

A networked robotic system reasons on a state-space representation that is exclusively discrete. This is possible when the continuous robot dynamics and the communication dynamics are neglible, i.e. operating with high reliability. For instance, the KIVA warehouse robots (Fig. 3.4., middle) are coordinated on a grid-like environment where robots move from center to center on a chessboard-like arrangement. Using markers on the floor, the robots are able to localize with high accuracy in a global coordinate system, which gives them the ability to move from cell to cell with a high reliability. By this the multi-robot system reduces to a centralized planning system on a grid, which can be analyzed with classical tools from algorithm analysis [Cormen *et al.* (2009)]. For instance the shortest path for a robot from one cell to the other could be calculated using Dijkstra's algorithm, and for N robots, collision free paths can be calculated in R^{2N} space [Everett *et al.* (1994); Parsons and Canny (1990)]. Due to the high-level of abstraction in deliberative planning systems, the multi-robot coordination problem becomes accessible to a wide range of methods from distributed algorithms that can provide provable correct solutions to resource allocation, consensus among distributed processes, data consistency, deadlock detection, leader election, among others [Lynch (1996)].

Deliberative algorithms fail, however, when the underlying continuous dynamics prove unreliable, and often times loose their analytical tractability when improved by backup mechanisms that take possible robot failure into account (see also [Parker (1998)] and Chap. **??** of this book). For instance in [Rutishauser *et al.* (2009)] a team of networked robots coordinates wirelessly to cover a grid-like environment. Robots have the ability to localize and exchange information on coverage progress. Using the Dijkstra algorithm, each robot calculates paths to the closest uncovered cell on the grid. As communication, navigation, and localization are unreliable the system has the following failure modes: coverage is redundant because robots fail to communicate, robots arrive late because they fail to navigate, and robots exchange false information due to errors in localization. The algorithm has therefore been adopted to maintain a probabilistic coverage map of the environment and have robots move towards the closest cell with the lowest likelihood of having been covered. [Rutishauser *et al.* (2009)] also shows how quantitative correct predictions of the system behavior can be obtained by calibrating the probability distributions of each failure mode.

As also the slightest uncertainty when executing a deliberative algorithm might break the system, verification of networked systems with stochastic subsystems (e.g. communication or actuation) is therefore a key challenge.

3.4 Challenges in Networked Robotic Systems

The science of networked robot systems lies at the intersection of mobility, control, and communication. Work focussing exclusively on control often has strong assumptions on the perception and communication abilities as well as their reliability. Recent advances in high-speed 6-DOF tracking, small holonomic aerial vehicles, and high-speed communication have enabled real-robot demonstrations of powerful formal approaches for formation control, collective motion and coverage, albeit relying on considerable infrastructure to overcome the limitations of on-board sensing (e.g. [Michael *et al.* (2009)]). On the other end of the spectrum are algorithms that are designed to explicitly deal with sensor and actuator noise by using random approaches (e.g. [Martinoli *et al.* (2004)]) or deliberative approaches that gracefully degrade to a randomized solution under influence of noise (e.g. [Amstutz *et al.* (2009)]). In order to bring these two orthogonal approaches together, we require for one methods to provide formal, probabilistic guarantees that allow to express system performance as a function of sensor and actuator reliability, and for another robust feedback controllers which stability is expressed as a function of sensor accuracy, actuator reliability, and communication rate. See also Section **??** of this book.

A promising application for networked robot systems is sensor coverage control, i.e. exploiting mobility for increasing the resolution of a sensor network. This field is maturing rapidly and solutions that are provably convergent and adaptive have been demonstrated on real robotic systems (e.g. [Schwager *et al.* (2009b)]). Few research has focussed on networked robot systems that can actively modify the environment based on sensor input. An instance of such a system is the distributed robotic garden where robots respond to watering requests of plants nested in sensing and communicating pots and thus create a feedback loop with the environment. In the future, such systems will need to combine identification of the underlying environmental model based on sensing with distributed control that can drive the system in a specific desired state.

Accurate mutual localization remains a key challenge. Current solutions relying on sound and light are susceptible to disturbance to cross-sensitivity and require substantial hardware investments, whereas radio-based and GPS-based systems lack the accuracy and bandwidth that allows for tight coordination. Promising directions are beam-forming MIMO antennas and small signal analysis. Before beam-forming antennas are wide-spread available, however, currently available systems that require mechanical rotation of a narrow beam antenna are bulky and slow.

Bibliography

(2009). Characterization and modeling of wireless channels for networked robotic and control systems — a comprehensive overview, in *Proc. of the Int. Conf. on Intelligent Robots and Systems (IROS)* (St. Louis, MO).

Agassounon, W., Martinoli, A. and Easton, K. (2004). Macroscopic modeling of aggregation experiments using embodied agents in teams of constant and time-varying sizes, *Auton Robot* **17**, 2-3, pp. 163–192.

Amstutz, P., Correll, N. and Martinoli, A. (2009). Distributed boundary coverage with a team of networked miniature robots using a robust market-based algorithm, *Annals of Mathematics and Artifcial Intelligence. Special Issue on Coverage, Exploration, and Search* **52**, 2–4, pp. 307–333.

Arkin, R. (1998). *Behavior-Based Robotics* (The MIT Press).

Bahl, P. and Padmanabhan, V. (2000). RADAR: An in-building RF-based user location and tracking system. in *IEEE Infocom*, Vol. 2, pp. 775–784.

Bhadauria, D. and Isler, V. (2009). Data gathering tours for mobile robots, in *Proceedings of the International Conference on Intelligent Robots and Systems (IROS)* (St. Louis, MO, USA).

Brooks, R. A. (1986). A robust layered control system for a mobile robot, *IEEE Journal of Robotics and Automation* **RA-2**, 1, pp. 14–23.

Caprari, G. and Siegwart, R. (2005). Mobile micro-robots ready to use: Alice, in *Proc. of the 2005 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2005)*, pp. 3295– 3300.

Chaimowicz, L., Michael, N. and Kumar, V. (2005). Controlling swarms of robots using interpolated implicit functions, *International Conference on Robotics* , pp. 2487 – 2492doi: 10.1109/ROBOT.2005.1570486.

Chen, Y. and Kobayashi, H. (2002). Signal strength-based indoor geolocation, in *IEEE International Conference on Communicatins*, Vol. 1, pp. 436–439.

Choi, J., Lee, J., Xu, Y. and Oh, S. (2010). Navigation strategies for swarm intelligence using spatio-temproal gaussian processes, *Special issue on Swarm Robotics, Neural Computation and Applications* To appear.

Cormen, T., Leiserson, C., Rivest, R. and Stein, C. (2009). *Introduction to Algorithms, Third Edition* (MIT Press).

Correll, N. (2007). *Coordination schemes for distributed boundary coverage with a swarm of miniature robots: synthesis, analysis and experimental validation*, Ph.D. thesis, Ecole Polytechnique Federale de Lausanne (EPFL).

Correll, N. (2008). Parameter estimation and optimal control of swarm-robotic systems: A case study in distributed task allocation, *Proc. of the 2008 IEEE Int. Conf. on Robotics and Automation (ICRA 2008)* , pp. 3302 – 3307.

Correll, N., Arechiga, N., Bolger, A., Bollini, M., Charrow, B., Clayton, A., Dominguez, F., Donahue, K., Dyar, S., Johnson, L., Liu, H., Patrikalakis, A., Robertson, T., Smith, J., Soltero, D., Tanner, M., White, L. and Rus, D. (2009a). Building a distributed robot garden, in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (St. Louis, MO), pp. 1509—1516.

Correll, N. and Martinoli, A. (2006a). Collective inspection of regular structures using a swarm of miniature robots, in M. H. A. Jr. and O. Khatib (eds.), *Proc. of the Ninth Int. Symp. on Experimental Robotics (ISER), Springer Tracts in Advanced Robotics*, Vol. 21 (Singapore), pp. 375–385.

Correll, N. and Martinoli, A. (2006b). System identification of self-organized robotic swarms, in *8th International Symposium on Distributed Autonomous Robotic Systems (DARS)*.

Correll, N. and Martinoli, A. (2006c). Towards optimal control of self-organized robotic inspection systems, in *In 8th International IFAC Symposium on Robot Control (SYROCO)* (Bologna, Italy).

Correll, N. and Martinoli, A. (2007). Modeling self-organized aggregation in a swarm of miniature robots, in *Proc. of the IEEE 2007 International Conference on Robotics and Automation Workshop on Collective Behaviors inspired by Biological and Biochemical Systems*.

Correll, N. and Martinoli, A. (2009). Multirobot inspection of industrial machinery, *IEEE Robotics & Automation Mag.* **16**, 1, pp. 103 – 112.

Correll, N., Rus, D., Bachrach, J. and Vickery, D. (2009b). Ad-hoc wireless network coverage with networked robots that cannot localize, in *IEEE International Conference on Robotics and Automation* (Kobe, Japan), p. 8 pages.

Dantu, K., Goyal, P. and Sukhatme, G. S. (2009). Relative bearing estimation from commodity radios, in *IEEE International Conference on Robotics and Automation* (IEEE), pp. 3871–3877.

Doniec, M., Vasilescu, I., Detweiler, C., Rus, D., Chitre, M. and Hoffmann-Kuhnt, M. (2009). Aquaoptical: A lightweight device for high-rate long-range underwater point-to-point communication, in *In Proceedings of OCEANS*, pp. papernumber 090601–157.

Estrin, D., Govindan, R., Heidemann, J. and Kumar, S. (1999). Next century challenges: scalable coordination in sensor networks, in *Proceedings of the 5th annual ACM/IEEE international conference on Mobile computing and networking*, pp. 263–270.

Everett, H. R., Gage, D. W., Gilbreath, G. A., Laird, R. T. and Smurlo, R. P. (1994). Real-world issues in warehouse navigation, in *Proceedings of the SPIE Conference on Mobile Robots IX*, Vol. 2352 (Boston, MA), pp. 629–634.

Frew, E. W. and Brown, T. X. (2008). Airborne communication networks for small unmanned aircraft systems, *Proceedings of the IEEE, Special Issue on Aviation Information Systems* **96**, 12, pp. 2008–2027.

Gurdan, D., Stumpf, J., Achtelik, M., Doth, K., Hirzinger, G. and Rus, D. (2006). Energy-efficient autonomous four-rotor flying robot controlled at 1khz, in *The 2006 International Conference on Robotics and Automation.*

Gutierrez, A., Campo, A., Dorigo, M., Amor, D., Magdalena, L., Monasterio-Huelin, F. and Madrid, S. (2008). An open localisation and local communication embodied sensor, *Sensors* **8**, pp. 7545–7563.

Hauert, S., Zufferey, J.-C. and Floreano, D. (2009). Evolved swarming without positioning information: an application in aerial communication relay. *Autonomous Robots* **26**, pp. 21–32.

Howard, A., Matarić, M. J. and Sukhatme, G. S. (2003). Putting the 'i' in 'team': An ego-centric approach to cooperative localization, in *IEEE International Conference on Robotics and Automation* (Taipei, Taiwan), pp. 868–892, URL `http://cres.usc.edu/cgi-bin/print_pub_details.pl?pubid=24`.

Howard, A., Parker, L. and Sukhatme, G. (2006). Experiments with a large heterogeneous mobile robot team: Exploration, mapping, deployment and detection, *Int. J. Robotics Research* **25**, 5-6, pp. 431–448.

Hsieh, M., Cowley, A., Kumar, V. and Taylor, C. (2008a). Maintaining network connectivity and performance in robot teams, *Journal of Field Robotics* **26**, 1–2, pp. 111–131.

Hsieh, M., Kumar, V. and Chaimowicz, L. (2008b). Decentralized controllers for shape generation with robotic swarms, *Robotics* **26**, 5, pp. 691–701.

Julian, B., Schwager, M., Angermann, M. and Rus, D. (2009). A location-based algorithm for multi-hopping state estimations within a distributed robot team, in *Proceedings of the 7th International Conference onField and Service Robotics (FSR 09)* (Cambridge, MA).

Kloetzer, M. and Belta, C. (2010). Automatic deployment of distributed teams of robots from temporal logic motion specifications, *IEEE Transactions on Robotics* **26**, 1, pp. 48–61.

Kottege, N. and Zimmer, U. (2008). Acoustical methods for azimuth, range and heading estimation in underwater swarms, in *Proceedings of Acoustics '08, Palais des Congress, Paris, France, 29 June-4 July 2008* (Paris, France), p. 5.

Kusy, B., Sallai, J., G. Balogh, A. . L., andJ. Tolliver, V. P., DeNap, F. and Parang, M. (2007). Radio interferometrictracking of mobile wireless nodes, in *ACM MobiSys*, pp. 139–151.

Lerman, K., Martinoli, A. and Galstyan, A. (2005). *Lect Notes Comput Sc* **3342**, pp. 143–152.

Lindhé, M., Johansson, K. H. and Bicchi, A. (2007). An experimental study of exploiting multipathfading for robot communications, in *Proceedings of Robotics: Science and Systems*, p. 8 pages.

Lochmatter, T. and Martinoli, A. (2009). *Experimental Robotics 11*, Vol. 54, chap. Tracking Odor Plumes in a Laminar Wind Field with Bio-Inspired Algorithms, springer tracts in advanced robotics edn. (Springer Verlag), pp. 473–482.

Lynch, N. (1996). *Distributed Algorithms* (Morgan Kaufman).

Maroti, M., Volgyesi, P., Dora, S., Branislav, K., andA. Ledeczi, A. N., Balogh, G. and Molnar, K. (2005). Radio interferometric geolocation, in *Proceedings of the 3rd internationalconference on Embedded networked sensor systems ACM SenSys*, pp. 1–12.

Martinoli, A., Easton, K. and Agassounon, W. (2004). Modeling swarm robotic systems: A case study in collaborative distributed manipulation, *Int. J. Robot Res.* **23**, 4-5, pp. 415–436.

Martinoli, A., Ijspeert, A. and Mondada, F. (1999). Understanding collective aggregation mechanisms: From probabilistic modelling to experiments with real robots, *Robotics and Autonomous Systems* **29**, 1, pp. 51–63.

Michael, N., Fink, J. and Kumar, V. (2009). Cooperative manipulation and transportation with aerial robots, in *Proc. of Robotics: Science and Systems* (Seattle, WA).

Mondada, F., Bonani, M., Raemy, X., Pugh, J., Cianci, C., Klaptocz, A., Magnenat, S., Zufferey, J.-C., Floreano, D. and Martinoli, A. (2009). The e-puck, a robot designed for education in engineering, in *Proc. 9th Conference on Autonomous Robot Systems and Competitions*, pp. 59–65.

Moshtagh, N., Michael, N., Jadbabaie, A. and Daniilidis, K. (2009). Vision-based, distributed control laws for motion coordination of nonholonomic robots, *IEEE Transactions on Robotics* **25**, 4, pp. 851–860.

Nagpal, R., Shrobe, H. and Bachrach, J. (2003). Organizing a global coordinate system from local information on an ad hoc sensor network, in *International Conference on Information Processing in Sensor Networks (IPSN)*.

Parker, L. E. (1998). ALLIANCE: An architecture for fault-tolerant multi-robotcooperation, *IEEE Transactions on Robotics and Automation* **14**, 2, pp. 220–240.

Parsons, D. and Canny, J. (1990). A motion planner for multiple mobile robots, in *Proc. IEEE International Conference on Robotics and Automation*, Vol. 1, pp. 8–13.

Priyantha, N., Chakraborty, A. and Balakrishnan, H. (2000). The cricket location-support system, in *Proceedings of the 6th annual international conference on Mobile Computing and Networking*, pp. 32–43.

Rudnick, D. L., Davis, R. E., Eriksen, C. C., Fratantoni, D. M. and Perry, M. J. (2004). Underwater gliders for ocean research, *MarineTechnology Society Journal* **38**, 1, pp. 48–59.

Rutishauser, S., Correll, N. and Martinoli, A. (2009). Collaborative coverage using a swarm of networked miniature robots, *Robotics & Autonomous Systems* **57**, 5, pp. 517–525.

Schill, F. and Zimmer, U. (2007). Pruning local schedules for efficient swarm communication, in *In Proc. of the IEEE International Symposium on Underwater Technology (UT '07)* (Tokyo, Japan).

Schwager, M., Julian, B. and Rus, D. (2009a). Optimal coverage for multiple hovering robots with downward-facing cameras, in *In Proc. of the International Conference of Robotics and Automation* (Kobe, Japan).

Schwager, M., Rus, D. and Slotine, J. J. (2009b). Decentralized, adaptive coverage control for networked robots, *International Journal of Robotics Research* **28**, 3, pp. 357–375.

Seybold, J. S. (2005). *Chapter 7: Near Earth Propagation Models*, Vol. 1 (John Wiley & Sons), pp.

134–146.

Shiomi, M., Kanda, T., Glas, D., Satake, S., Ishiguro, H. and Hagita, N. (2009). Field trial of networked social robots in a shopping mall, in *Proceedings of the 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems* (St. Louis, MO).

Spears, W. M., Spears, D. F., Heil, R., Kerr, W. and Hettiarachchi, S. (2005). An overview of physicomimetics, in *Swarm Robotics*, *Lecture Notes in Computer Science*, Vol. 3342 (Springer, Berlin), pp. 84–97.

Stomp, M., Huisman, J., Stal, L. and Matthijs, H. (2007). Colorful niches of phototrophic microorganisms shaped by vibrations of the water molecule, *The ISME Journal: Multidisciplinary Journal of Microbial Ecology* **1**, pp. 271–282.

Tanenbaum, A. (2002). *Computer Networks* (Prentice Hall).

Vasilescu, I., Varhavskaya, P., Kotay, K. and Rus, D. (2005). Autonomousmodular optical underwater robot (amour) design, prototype and feasibility study, in *In Proceedings of the International Conference on Robotics and Automation (ICRA)* (Barcelona, Spain), pp. 1603–1609.

Winfield, A., Liu, W., Nembrini, J. and Martinoli, A. (2008). Modelling a wireless connected swarm of mobile robots, *Swarm Intell* **2**, 2, pp. 241–266.

Yun, S., Schwager, M. and Rus, D. (2009). Coordinating construction of truss structures using distributed equal-mass partitioning, in *Proc. of the 14th International Symposium on Robotics Research* (Luzern, Switzerland).